

Equilibrium Refinements

Mihai Manea

MIT

Sequential Equilibrium

- ▶ In many games information is imperfect and the only subgame is the original game. . . subgame perfect equilibrium = Nash equilibrium
- ▶ Play starting at an information set can be analyzed as a separate subgame if we specify players' **beliefs** about at which node they are.
- ▶ Based on the beliefs, we can test whether continuation strategies form a Nash equilibrium.
- ▶ **Sequential equilibrium** (Kreps and Wilson 1982): way to derive plausible beliefs at every information set.

An Example with Incomplete Information

Spence's (1973) job market signaling game

- ▶ The worker knows her ability (productivity) and chooses a level of education.
- ▶ Education is more costly for low ability types.
- ▶ Firm observes the worker's education, but not her ability.
- ▶ The firm decides what wage to offer her.

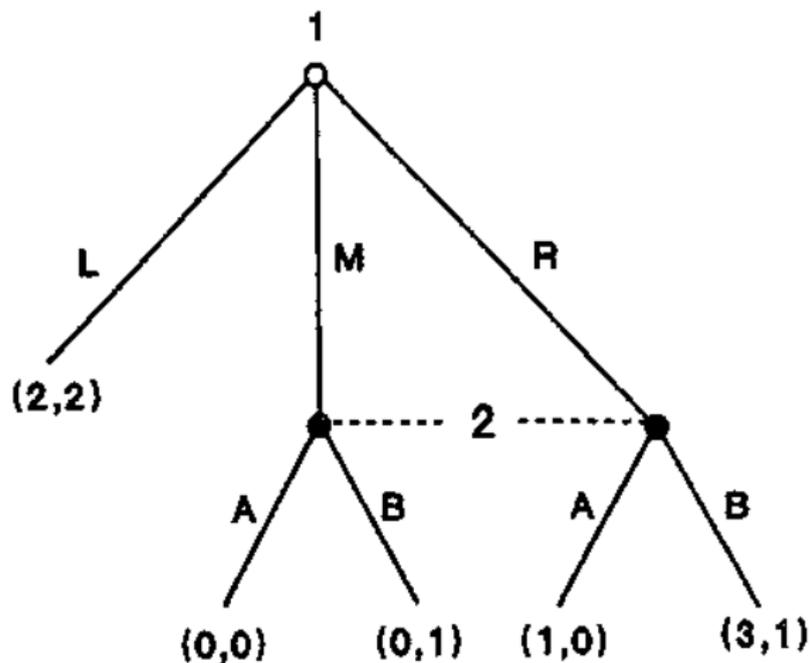
In the spirit of subgame perfection, the optimal wage should depend on the firm's beliefs about the worker's ability given the observed education.

An equilibrium needs to specify contingent actions and beliefs.

Beliefs should follow Bayes' rule on the equilibrium path.

What about off-path beliefs?

An Example with Imperfect Information



7cl fH'gmcZHNY'A -H'DfYgg"l gYX k Jh' dYfa Jggjcb"

Figure: (L, A) is a subgame perfect equilibrium. Is it plausible that 2 plays A?

Assessments and Sequential Rationality

Focus on extensive-form games of perfect recall with finitely many nodes.

An **assessment** is a pair (σ, μ)

- ▶ σ : (behavior) strategy profile
- ▶ $\mu = (\mu(h) \in \Delta(h))_{h \in H}$: **system of beliefs**

$u_i(\sigma|h, \mu(h))$: i 's payoff when play begins at a node in h randomly selected according to $\mu(h)$, and subsequent play specified by σ .

The assessment (σ, μ) is **sequentially rational** if

$$u_{i(h)}(\sigma_{i(h)}, \sigma_{-i(h)}|h, \mu(h)) \geq u_{i(h)}(\sigma'_{i(h)}, \sigma_{-i(h)}|h, \mu(h))$$

for all information sets h and alternative strategies σ' .

Consistency

Beliefs need to be consistent with strategies.

$\tilde{\sigma}$ is totally mixed if $\text{supp}(\tilde{\sigma}_{i(h)}(h)) = A(h)$, i.e., all information sets are reached with positive probability.

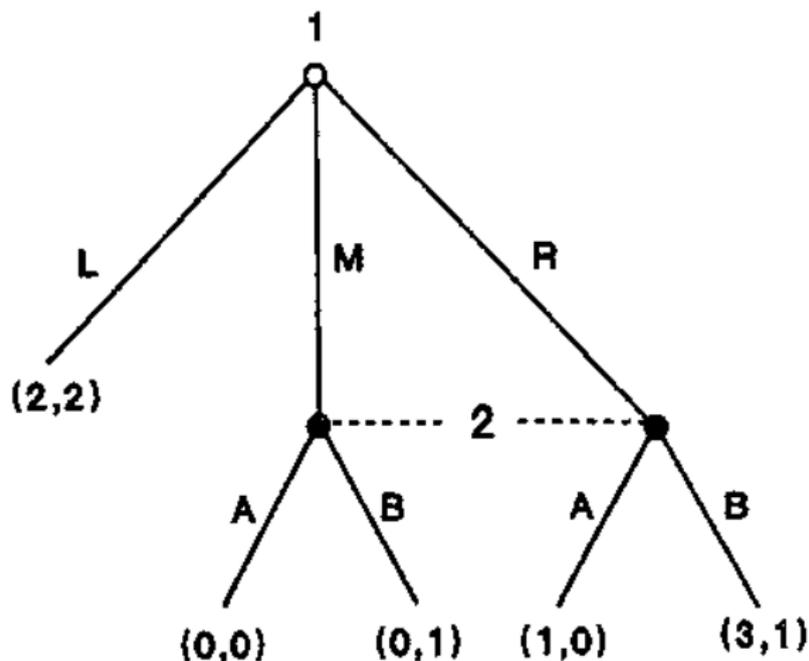
Bayes' rule \rightarrow unique system of beliefs $\mu^{\tilde{\sigma}}$ for any totally mixed $\tilde{\sigma}$.

The assessment (σ, μ) is **consistent** if there exists a sequence of totally mixed strategy profiles $(\sigma^m)_{m \geq 0} \rightarrow \sigma$ s.t. $(\mu^{\sigma^m})_{m \geq 0} \rightarrow \mu$.

Definition 1

A *sequential equilibrium* is an assessment that is sequentially rational and consistent.

Implications of Sequential Rationality



7ci fH#gmcZH\Y A +H DFYgg*1 gYX k J\k dYfa jggjcb*

Figure: No belief rationalizes A. 2 plays B, 1 optimally chooses R.

Implications of Consistency

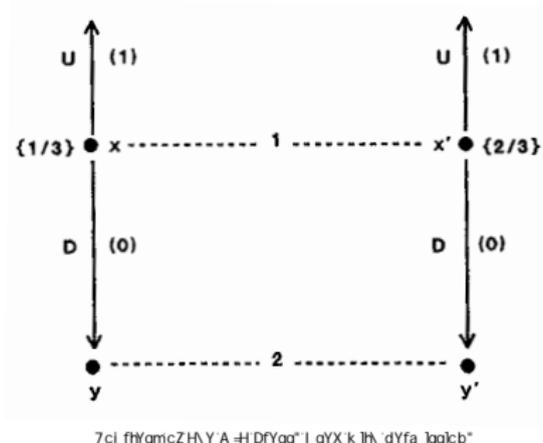


Figure: By consistency, $\mu(y|h_2) = \mu(x|h_1)$, even though D is never played.

Consistency \rightarrow **common beliefs** after deviations from equilibrium behavior.

Why should different players have the same theory about something not supposed to happen?

Consistency matches the spirit of equilibrium analysis, which assumes players hold identical beliefs about others' strategies.

Existence of Sequential Equilibrium

Theorem 1

A sequential equilibrium exists for every finite extensive-form game.

Follows from existence of perfect equilibria, prove later.

Sequential Equilibrium Multiplicity

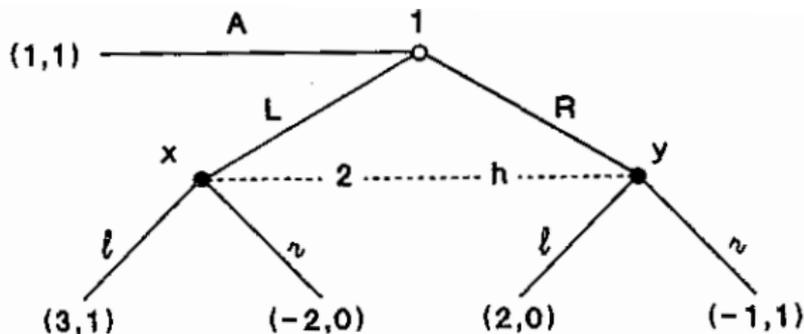
Theorem 2

*For generic payoff functions, the set of sequential equilibrium **outcome distributions** is finite.*

Set of sequential equilibrium **assessments** often infinite

- ▶ Infinitely many belief specifications at off-path information sets supporting some equilibrium strategies.
- ▶ Set of sequential equilibrium strategies may also be infinite. Off-path information sets may allow for consistent beliefs that make players indifferent between actions. . . many mixed strategies compatible with sequential rationality.

Example



7cl fHrgmcZHVY'A +H DFYgg"l gYX k jh' dYfa Jggjcb"

Sequential equilibrium outcomes: (L, l) and A

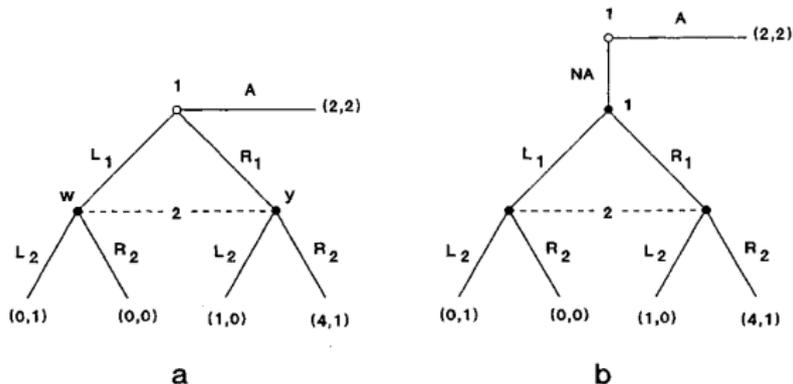
Unique equilibrium leading to (L, l)

Two families of equilibria with outcome A... 2 must choose r with positive probability

- 1 2 chooses r with probability 1 and believes $\mu(x) \in [0, 1/2]$
- 2 2 chooses r with probability in $[2/5, 1]$ and believes $\mu(x) = 1/2$

Sequential Equilibrium Is Sensitive to the Extensive Form

“Strategically neutral” changes in game tree affect equilibria.



7ci fH/gmcZH.Y A -H DFYgg* I gYX k Jh dYfa jggjcb*

Game a: (A, L_2) possible in a sequential equilibrium

Game b: $((NA, R_1), R_2)$ unique sequential equilibrium strategies. In subgame following NA , R_1 strictly dominates L_1 . Then 2 chooses R_2 , and 1 best responds with (NA, R_1) .

Perfect Equilibrium

	L	R
U	1,1	0,0
D	0,0	0,0

Selten (1975): (trembling-hand) perfect equilibrium

- ▶ Both (U, L) and (D, R) are Nash equilibria.
- ▶ (D, R) not robust to small **mistakes**: if 1 thinks that 2 might make a mistake and play L with positive probability, deviate to U .

Definition 2

In a strategic-form game, a profile σ is a *perfect equilibrium* if there is a sequence of *trembles* $(\sigma^m)_{m \geq 0} \rightarrow \sigma$, where each σ^m is a totally mixed strategy, such that σ_i is a best reply to σ_{-i}^m for each m and all $i \in N$.

Existence of Perfect Equilibria

Definition 3

σ^ε is an ε -perfect equilibrium if $\exists \varepsilon(s_i) \in (0, \varepsilon], \forall i \in N, s_i \in S_i$ s.t. σ^ε is a Nash equilibrium of the game where players are restricted to play mixed strategies in which every pure strategy s_i has probability at least $\varepsilon(s_i)$.

Proposition 1

A strategy profile is a perfect equilibrium iff it is the limit of a sequence of ε -perfect equilibria as $\varepsilon \rightarrow 0$.

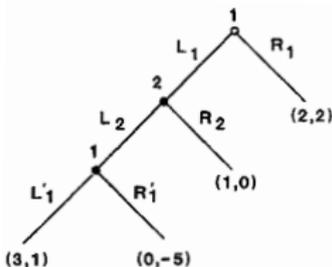
Theorem 3

Every finite strategic-form game has a perfect equilibrium.

Proof.

A $1/n$ -perfect equilibrium exists by the general Nash equilibrium existence theorem. By compactness, the sequence of $1/n$ -perfect equilibria has a convergent subsequence as $n \rightarrow \infty$. The limit is a perfect equilibrium. \square

Perfection in Strategic Form \Rightarrow Subgame-Perfection



a. Extensive Form

	L_2	R_2
R_1	2,2	2,2
L_1, L'_1	3,1	1,0
L_1, R'_1	0,-5	1,0

b. Strategic Form

7ci fH/gmcZ H.V A -H DfYgg" l gYX k]h\ dYfa]gg]cb*

Unique SPE: $(L_1 L'_1, L_2)$

(R_1, R_2) is perfect in strategic form, sustained by trembles s.t. after trembling to L_1 , player 1 chooses R'_1 vs. L'_1 with probability ratio $\geq 1/5$.

Correlation in trembles at different information sets... unreasonable.

Perfection in Extensive-Form Games

Solution: **agent-normal form**

- ▶ A different player for every information set h .
- ▶ “Player” h has the same payoffs as $i(h)$.

Definition 4

A *perfect equilibrium* for an extensive-form game is a perfect equilibrium of its agent-normal form.

Connection to Sequential Equilibrium

Theorem 4

Every perfect equilibrium of a finite extensive-form game is a sequential equilibrium (for some appropriately chosen beliefs).

- ▶ σ : perfect equilibrium of the extensive-form game $\Rightarrow \exists (\sigma^m)_{m \geq 0} \rightarrow \sigma$ totally mixed strategies in the agent-normal form s.t. σ_h is a best reply to σ_{-h}^m for each m and all information sets h .
- ▶ By compactness, $(\mu^{\sigma^m})_{m \geq 0}$ has a convergent subsequence, denote limit by μ .
- ▶ By construction, (σ, μ) is consistent.
- ▶ σ_h is a best response to $\mu^{\sigma^m}(h)$ and σ_{-h}^m for each m .
- ▶ By continuity, σ_h is a best response to $\mu(h)$ and σ_{-h} .
- ▶ One-shot deviation principle: (σ, μ) is sequentially rational

Proper Equilibrium

Myerson (1978): a player is infinitely more likely to tremble to better actions

A player's probability of playing the second-best action is at most ε times the probability of the best, the probability of the third-best action is at most ε times the probability of the second-best. . .

Definition 5

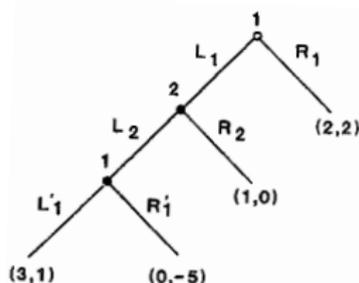
An ε -proper equilibrium is a totally mixed strategy profile σ^ε s.t. if $u_i(s_i, \sigma_{-i}^\varepsilon) < u_i(s'_i, \sigma_{-i}^\varepsilon)$, then $\sigma_i^\varepsilon(s_i) \leq \varepsilon \sigma_i^\varepsilon(s'_i)$. A proper equilibrium is any limit of ε -proper equilibria as $\varepsilon \rightarrow 0$.

Theorem 5

Every finite strategic-form game has a proper equilibrium.

Prove existence of ε -proper equilibria applying Kakutani's fixed point theorem to "mistake hierarchy ε -best response" correspondences, then use compactness to find a limit point.

Properness in Strategic Form \Rightarrow Subgame Perfection



a. Extensive Form

	L ₂	R ₂
R ₁	2,2	2,2
L ₁ ,L' ₁	3,1	1,0
L ₁ ,R' ₁	0,-5	1,0

b. Strategic Form

7ci ffit/gmcZHNY A +H DFYgg' l gYX k Jk dYfa jggjcb*

Forward Induction

Equilibrium: off-path observations interpreted as errors

Forward induction: players should believe in the rationality of their opponents even after observing deviations.

- ▶ When a player deviates from equilibrium strategies, the opponent should believe that the player expects follow up play that makes the deviation reasonable.
- ▶ The deviation is informative about the player's type or, in general extensive form games, about his future play.

Forward induction not an equilibrium concept: in equilibrium, all players expect specified strategies to be exactly followed

An attempt to describe strategic uncertainty. . . no single, rigorous definition

Example

1 chooses between O , which generates payoffs $(2, 2)$, or I , which leads to

	T	W
T	0,0	3,1
W	1,3	0,0

SPE: (OW, T) . Reasonable?

- ▶ If 1 plays I , this suggests he does not intend to follow up with W : O yields a payoff of 2, while W leads to a payoff of at most 1 for player 1.
- ▶ Player 2, anticipating that 1 will play T , should play W .
- ▶ If 1 can convince 2 to play W , he gets the higher payoff from (T, W) .

Forward Induction and Strict Dominance

Reduced normal form

	T	W
O	2,2	2,2
IT	0,0	3,1
IW	1,3	0,0

(O, T) is a perfect (in fact, proper) equilibrium.

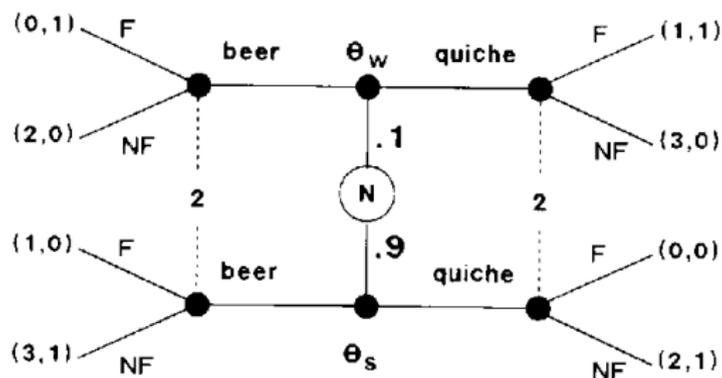
If we rule out IW because it is s. dominated by O , then the only perfect equilibrium is (IT, W) .

Signaling Games

- ▶ Two players: sender S and receiver R
- ▶ T : set of types for S
- ▶ $p(t)$: probability of type $t \in T$
- ▶ S privately observes his type t , then sends a message $m \in M(t)$
- ▶ $T(m) = \{t \mid m \in M(t)\}$: types that can send message m
- ▶ R observes m and chooses an action $a \in A(m)$
- ▶ Payoffs $u_S(t, m, a)$ and $u_R(t, m, a)$

If S plays m with probability 0, then any beliefs for R about t after observing m are consistent. . . sequential equilibrium imposes no restrictions on beliefs off the equilibrium path in signaling games.

The Beer-Quiche Game



7ci fh/gmcZHX' A +H DFYgg" l gYX k Jh' dYfa jggjcb"

- ▶ Player 1 is wimpy (w) or surly (s), with probabilities .1 and .9; $T = \{w, s\}$.
- ▶ 1 orders breakfast: $M = M(t) = \{beer, quiche\}, \forall t \in T$.
- ▶ Player 2 decides whether to fight: $A(m) = \{F, NF\}, \forall m \in M$.
- ▶ 1 gets utility 1 from having his favorite breakfast—beer if surly, quiche if wimp—but a disutility of 2 from fighting.
- ▶ When 1 is w , 2's payoff is 1 if he fights and 0 otherwise; when 1 is s , payoffs are reversed.

Sequential Equilibria

All sequential equilibria involve **pooling**

- ▶ Compare $\sigma_2(F|beer)$ and $\sigma_2(F|quiche)$
- ▶ Breakfast leading to a smaller probability of fighting must be selected with probability 1 in equilibrium by player 1 type who likes it. . .

Classes of sequential equilibria

- 1 Both types of player 1 drink beer.
- 2 Both types of player 1 eat quiche.

Player 2 does not fight in equilibrium. Player 2 must fight with probability at least 1/2 when observing the out-of-equilibrium breakfast. . . supported by any belief for player 2 placing probability at least 1/2 on w following the out-of-equilibrium breakfast.

Forward-Induction in Beer-Quiche

Quiche equilibrium unreasonable

- ▶ Why would the wimp deviate to beer? No matter how 2 reacts, wimp cannot get more than 2, and he is already getting 3.
- ▶ Seeing beer, 2 should conclude that 1 is surly and not fight, which would induce surly type to deviate.

Forward-induction argument does not rule out the beer equilibrium

- ▶ In the beer equilibrium, it is unreasonable for surly type to deviate to quiche, while reasonable for wimp.
- ▶ 2's belief that 1 is wimpy if he orders quiche is reasonable.

Intuitive Criterion

Cho and Kreps (1987): intuitive criterion

- ▶ Robustness to replacing the equilibrium path by its expected payoff
- ▶ Presumes that players are certain about play on the equilibrium path, but there is uncertainty off the path
- ▶ If m can never lead to a higher payoff for t than his equilibrium payoff, then equilibrium beliefs should assign probability 0 to type t following m .

Irrational Strategies for the Receiver

What if 2 can also pay a milion dollars to 1?

- ▶ It would be reasonable for both types to deviate.
- ▶ But 2 would never want to pay a million dollars.
- ▶ Assume 1 cannot expect 2 to play a irrational strategy.

Intuitive Criterion

For any $T' \subseteq T$ and any message m ,

$$BR(T', m) = \cup_{\mu \mid \mu(T')=1} BR(\mu, m)$$

for strategies that R could rationally play after m and if he is certain that $t \in T'$.

Consider a sequential equilibrium

- ▶ $u_S^*(t)$: equilibrium payoff to type t
- ▶ $\tilde{T}(m) = \{t \mid u_S^*(t) > \max_{a \in BR(T(m), m)} u_S(t, m, a)\}$: types that do better in equilibrium than they could possibly do by sending m , no matter how R reacts, as long as R is rational.

The equilibrium **fails the intuitive criterion** if $\exists t' \in T, m \in M(t')$ s.t.

$$u_S^*(t') < \min_{a \in BR(T(m) \setminus \tilde{T}(m), m)} u_S(t', m, a).$$

Discussion

The equilibrium fails the intuitive criterion if some sender type is getting less than any payoff he could possibly get by playing m , assuming he could convince the sender that he is not in $\tilde{T}(m)$ because m does not make sense for any of those types.

In the beer-quiche example, the quiche equilibrium fails this criterion.

Spence's Signaling Mode

*The following slides are based on lecture notes by Debraj Ray.

Spence's (1973) job market signaling game

- 1 Nature chooses a worker type (ability) $\theta \in \{H, L\}$ with $H > L > 0$; the probability of H is $p \in (0, 1)$.
- 2 Type is revealed to the worker but not to the employer (firm).
- 3 Worker chooses $e \geq 0$ units of education, incurs disutility e/θ .
- 4 Firm observes e , forms an estimate of θ , and pays the worker wage $\mathbb{E}(\theta|e)$ (perfect competition: firm has 0 expected payoff).

Payoff of type θ worker: $\mathbb{E}(\theta|e) - e/\theta$. $\mathbb{E}(\theta|e)$ endogenously derived from strategies. . . depends on worker strategies and how they translate into beliefs.

If worker chooses e with probability 0, then any belief about θ after observing e is consistent. . . sequential equilibrium imposes no restrictions on off-path beliefs.

Single Crossing

Lemma 1

If H and L choose e and e' , respectively, with positive probability in equilibrium, then $e \geq e'$.

- ▶ H does not have incentives to deviate from e to e' ,

$$\mathbb{E}(\theta|e) - \frac{e}{H} \geq \mathbb{E}(\theta|e') - \frac{e'}{H}.$$

- ▶ L does not have incentives to deviate from e' to e ,

$$\mathbb{E}(\theta|e') - \frac{e'}{L} \geq \mathbb{E}(\theta|e) - \frac{e}{L}.$$

- ▶ Adding the two inequalities,

$$(e - e') \left(\frac{1}{L} - \frac{1}{H} \right) \geq 0.$$

Key assumption: higher types have lower marginal cost (result holds for cost functions other than e/θ).

Separating Equilibrium

Each type takes a different action (type perfectly revealed)

- ▶ L must choose $e = 0$, equilibrium wage L for equilibrium action
- ▶ H cannot mix, profitable deviation to lowest action in support

Incentive constraints if H chooses e^*

- ▶ L does not want to imitate H

$$L \geq H - \frac{e^*}{L} \Rightarrow e^* \geq L(H - L) =: e_1$$

- ▶ H does not want to imitate L

$$H - \frac{e^*}{H} \geq L \Rightarrow e^* \leq H(H - L) =: e_2$$

Any $e^* \in [e_1, e_2]$ is possible in a sequential equilibrium with suitably chosen off-path beliefs. E.g., employer believes that any $e < e^*$ comes from L , while $e > e^*$ from H .

Pooling Equilibrium

Both types play the same action e^* (no information is revealed)

- ▶ Equilibrium wage: $pH + (1-p)L$
- ▶ Off-path wage? Depends on off-path beliefs.
- ▶ Strongest incentives to follow equilibrium play: if firm believes worker is type L for $e \neq e^*$, then wage should be L .

Neither type should want to deviate to 0,

$$pH + (1 - p)L - \frac{e^*}{\theta} \geq L, \quad \forall \theta \in \{H, L\}.$$

Binding constraint for $\theta = L$,

$$e^* \leq pL(H - L).$$

Hybrid Equilibria

One or both types mix (partial info revelation). An example

- ▶ L chooses 0
- ▶ H chooses 0 with probability q and e with probability $1 - q$ for some $q \in (0, 1)$ and $e > 0$
- ▶ After observing e , firm believes worker type is H , offers wage H .
- ▶ After 0, worker type is H with probability $qp/(qp + 1 - p)$, offers wage

$$\frac{qp}{qp + 1 - p}H + \frac{1 - p}{qp + 1 - p}L.$$

- ▶ H must be indifferent between 0 and e (then L does not have incentives to deviate to e),

$$\frac{qp}{qp + 1 - p}H + \frac{1 - p}{qp + 1 - p}L = H - \frac{e}{H}.$$

- ▶ Off-path beliefs and wages as before

Intuitive Criterion

In Spence's signaling model, all three types of equilibria—separating, pooling, and hybrid equilibria—coexist.

Beliefs freely assigned off the equilibrium path (consistency has no bite).

Apply the intuitive criterion to get sharper predictions.

Proposition 2

A single equilibrium outcome survives the intuitive criterion—the separating equilibrium in which L plays 0 while H plays $e_1 = L(H - L)$.

This is the most efficient separating equilibrium.

Proof

First rule out pooling and hybrid equilibria. Suppose both H and L play e with positive probability.

- ▶ λ : firm's posterior belief that worker is type H after e
- ▶ payoff of type θ after choosing e : $\lambda H + (1 - \lambda)L - e/\theta$

Let $e' > e$ solve

$$H - \frac{e'}{L} = \lambda H + (1 - \lambda)L - \frac{e}{L}.$$

Choose $e'' > e'$ close to e' . The equilibrium fails the intuitive criterion for type H and message e'' .

- ▶ L would not deviate to e'' even if firm offers wage H after e'' ,

$$\lambda H + (1 - \lambda)L - \frac{e}{L} > H - \frac{e''}{L}.$$

- ▶ H would deviate to e'' if firm was convinced of H after e'' ,

$$\lambda H + (1 - \lambda)L - \frac{e}{H} < H - \frac{e''}{H}.$$

Proof

Rule out other separating equilibria. Consider a separating equilibrium where L plays 0 and H plays $e > e_1$. Fix $e' \in (e_1, e)$. The equilibrium fails the intuitive criterion for type H and message e' .

- ▶ L would not deviate to e' even if firm believes worker is type H after e' ,

$$H - \frac{e'}{L} < H - \frac{e_1}{L} = L.$$

- ▶ H would deviate to e' if that convinces firm that worker is type H ,

$$H - \frac{e}{L} > H - \frac{e'}{L}.$$

MIT OpenCourseWare

<https://ocw.mit.edu>

14.16 Strategy and Information

Spring 2016

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.